AD-A200 938

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

MACHINE INTELLIGENCE APPLIED TO RADAR OBJECT MODELING

A.M. AULL
R.A. GABEL
Group 21

TECHNICAL REPORT 818

12 OCTOBER 1988

LEXINGTON MASSACHUSETTS

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This program is sponsored by the U.S. Army Strategic Defense Command under Air Force Contract F19628-85-C-0002.

The views and conclusions contained in this document are those of the contractor and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the United States Government.

The ESD Public Affairs Office has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Hugh L. Southall

Hugh L. Southall, Lt. Col., USAF
Chief, ESD Lincoln Laboratory Project Office

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission is given to destroy this document when it is no longer needed.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

MACHINE INTELLIGENCE
APPLIED TO RADAR OBJECT MODELING

A.M. AULL
R.A. GABEL
Group 21

TECHNICAL REPORT 818

12 OCTOBER 1988

LEXINGTON MASSACHUSETTS

ABSTRACT

A point of commonality among sensor data understanding problems is the inability of any single classical approach to illuminate clearly and robustly the amount of information that is inherent to the data and is usually apparent to a human observer. Therefore, there is appreciable synergism among various sensor data understanding approaches as they draw from the combined fields of signal processing, image processing, pattern recognition, estimation and decision theory, and, more recently, machine intelligence. Our particular application is in the area of range-Doppler images of simple space objects. These images are characterized by their relatively sparse detail, coupled with simple but rapid motion. The data is multidimensional in nature with additive noise, distortion, and missing points. The relevant features to be extracted from the radar data include the position of all scattering centers in body coordinates, identification of the scattering center type (sphere, corner, edge, etc.), and motion parameters for the object.

Our goal is to produce a representation of the imaged 3-dimensional object that is appropriate for recognizing the object as an example of something we have seen before and cataloged, recognizing the object as an uncataloged object, or determining discrepancies between the recognized object and our expectations of its appearance. The motivation in our approach is thus to combine the power of classical methods with the flexible representations and control structures afforded by the field of machine intelligence and to derive primitive features from which a semantic model can be built. We have found that the semantic model is readily manipulated and interpreted by a human observer as well as by recognition, discrimination, and generalization procedures.

We have built a recognition system with three major conceptual modules. The first of these is a set of signal processing primitives that are directed at the data to select subsets of data, extract features, and compare extracted features with the data to produce confidence measures. The second major module is the semantic model building and matching scheme. This component takes the data-derived features and produces a semantic model which is then matched against a catalog of stored semantic models for object identification. The final conceptual module in the system is the control structure which is based on a blackboard architecture from the field of machine intelligence. In this report, we will discuss the issues in range-Doppler imaging and describe the major system modules in more detail including motivations for our approach. We will also provide some performance evaluation results. Finally, summary and directions for future work will be presented.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

1. INTRODUCTION

The study of sensor data understanding is motivated by a variety of goals including discrimination, object identification, and threat assessment. The extent to which the primary goal affects the technical direction is dependent upon the sensor physics and the nature of the observed objects. A point of commonality among sensor data understanding problems is the inability of any single classical approach to illuminate clearly and robustly the amount of information that is inherent to the data and is usually apparent to a human observer. Therefore, there is an appreciable synergism among various sensor data understanding approaches as they draw from the combined fields of signal processing, image processing, pattern recognition, estimation and decision theory, and, more recently, machine intelligence [1].

Our particular application is in the area of image understanding. More specifically, we are given a time sequence of range-Doppler images of simple space objects. These range-Doppler images are characterized by their relatively sparse detail, coupled with simple but rapid motion. The detail is on the order of 10 scattering centers per image, with object spin and precession at rates of a few cycles per second. The physical features of the vehicle typically present one scattering center apiece, compared with the more complex detail of, for example, a spacecraft solar panel. The data is multidimensional in nature with additive noise, distortion, and missing points. The interesting features to be extracted from the radar data include the position of all scattering centers in body coordinates, identification of the scattering center type (sphere, corner, edge, etc.), and motion parameters for the object. A source of complexity is that the features can be time-varying and change abruptly. Our goal is to produce a representation of the imaged 3-dimensional object that is appropriate for recognizing the object as an example of something we have seen before and cataloged, recognizing the object as an occurrence of an uncataloged object, or determining discrepancies between the recognized object and our expectations of its appearance.

Conventional approaches typically involve comparing derived and cataloged representations obtained from standard signal processing algorithms, parametric models, or geometric primitives. Although these approaches can be very powerful, they are also computationally intensive, sensitive to noise, and not representative of the signal in a form that is readily interpreted or edited. We have examined a variety of scattering center histories and found that the human observer can successfully extract an estimate of scattering center motion from even noisy data at low imaging rates. This reinforces the claim that a knowledge-based approach can more successfully model an observed object than can a more traditional signal processing approach.

The motivation in our approach is thus to combine the power of classical methods with the flexible representations and control structures afforded by the field of machine intelligence. A system goal is to derive primitive features from which a semantic model can be built. The semantic model is readily manipulated and interpreted by a human observer as well as by recognition, discrimination, and generalization procedures.

We have built a recognition system and interactive workstation which we call the Radar Object Modeling Environment or ROME. There are three major conceptual modules in our current

system. The first of these is a set of signal processing primitives that are directed at the data to select subsets of data, extract features, and compare extracted features with the data to produce confidence measures.

The second major module is the semantic model building and matching scheme. This component takes the data-derived features and produces a semantic model which is then matched against a catalog of stored semantic models for object identification. The semantic model represents the extracted information in terms of physical components, relationships between the components, and properties of the components and their relationships. The extracted model and the stored catalog model are in identical form and thus can readily be compared and contrasted. In addition, several examples of semantic models can be handed to the system which will generalize a single semantic model representation. This feature is useful when presented with data of objects for which we do not have *a priori* information.

The final conceptual module in the system is the control structure which is based on blackboard theory [2] from the field of machine intelligence. The blackboard structure consists of a global memory (the "blackboard") where information is posted as it is accumulated. Knowledge sources which embody specific information processing algorithms check the state of the blackboard to see if their expertise is needed. If a knowledge source is called on, it extracts information currently on the blackboard, does its own processing, and posts its results back to the blackboard. The control continues in this manner until there are no remaining candidate knowledge sources. A scheduler maintains control of triggering knowledge sources in cases when more than one is ready to perform some action. The blackboard design is inherently *opportunistic* and modular, and allows for both data-driven and model-driven processing simultaneously.

The user interface for ROME displays the recognition and blackboard control operations as they proceed. ROME's ability to exhibit the control actions facilitates the development of better control strategies within the blackboard framework. In addition to its recognition capabilities, ROME can act as a workstation for manual analysis of target characteristics. The menu-driven workstation provides a useful interactive analysis tool and allows users to learn from the manual operation how to build better deductive procedures.

In this report, we will discuss the issues in range-Doppler imaging and describe ROME's major system modules in more detail including motivations for our approach. We will also provide some limited performance evaluation results. Finally, summary and directions for future work will be presented.

2. RADAR RANGE-DOPPLER IMAGING

The data used in this analysis is taken from a range-Doppler imaging radar, which produces a sequence of two-dimensional portraits of the target being observed. Refer to Figure 2-1, which shows two scattering centers (i.e., points of reflection) on a target. Assume that the radar tracker

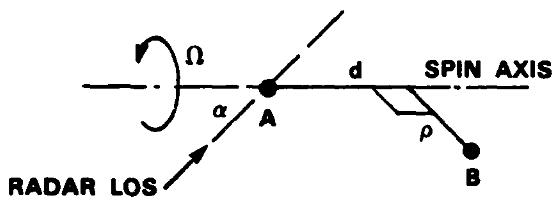

Figure 2-1. Simplified Target Model.

has removed the gross target motion, and that we measure the relative range of scatterer B with respect to scatterer A. Assume that the scatterers are part of a rigid body spinning about its axis with angular rate Ω radians per second, with B at axial distance d and radial distance ρ with respect to A, and with angle α between the spin axis and the radar line-of-sight (LOS). The range increment from A to B along the radar LOS is given by

$$r(t) = d\cos(\alpha) + \rho\sin(\alpha)\sin(\Omega t). \tag{2.1}$$

Suppose that a sequence of pulses is collected at a fixed pulse repetition interval. Evaluating the time-rate-of-change of this relative range, for example by performing an FFT of the time- and range-sampled data, we obtain a range rate of $v(t) = \rho\Omega\sin(\alpha)\cos(\Omega t)$ for point B with respect to point A, observed as a Doppler shift of

$$f_D(t) = \frac{2v(t)}{\lambda} = \frac{2\rho\Omega}{\lambda}\sin(\alpha)\cos(\Omega t). \tag{2.2}$$

Combining 2.1 and 2.2, we see that the locus of scatterer B forms an elliptical path in the range-Doppler plane, as shown in Figure 2-2. Since the Doppler frequency f_D is directly proportional to ρ, the Doppler dimension is also commonly referred to as the crossrange axis. Note that for a signal bandwidth of W Hertz and imaging interval of T seconds, with corresponding time and frequency resolutions $\delta t = 1/W$ and $\delta f_D = 1/T$, the range and crossrange resolutions

$$\delta r = \frac{c\,\delta t}{2} = \frac{c}{2W} \quad \text{and} \quad \delta\rho = \frac{\lambda\,\delta f_D}{2\Omega} = \frac{\lambda}{2\Omega T} \tag{2.3}$$

are independent of the target range. The only assumption is that the imaging interval T is smaller than the time taken by the imaged point to move from one range-Doppler cell to another. (For more details and a relaxation of this requirement, see [3].) Note how this presentation of data differs from the optical case — rather than the conventional angle-angle scan, we illuminate the

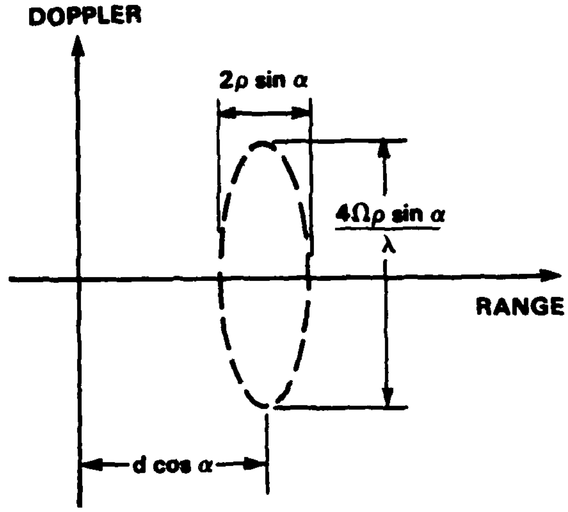

Figure 2-2. Scatterer Locus.

object from the radar direction and project the reflection intensity onto a plane under the object, parallel to the radar LOS. Note, too, that the assumed motion is necessary to form the image: with $\Omega = 0$, no crossrange information is obtained.

We can use this analysis to model a complex rigid body as a collection of scatterers whose projections in the range-Doppler plane traverse elliptical paths, perhaps with occlusion as individual scatterers are masked from the radar illumination by another part of the object. Looking at successive range-Doppler images, we will observe motion of the rotating scattering centers. This motion may be exhibited graphically as a plot of scattering center locations in range-Doppler-time space, as seen in Figure 2-3—a rather confusing collection of points, each of which is associated

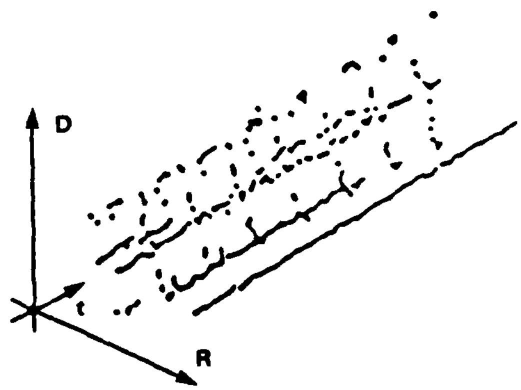

Figure 2-3. Scatterers in Range-Doppler-Time Space.

with an amplitude peak from a (complex-valued) image. This forms the input to our image understanding system: given data such as that presented in Figure 2-3, our task is to produce a model in the form of Figure 2-1 that explains this data. Our task is complicated by the occurrence of extra points due to noise-induced peaks in the range-Doppler plane, by missing data due to dropped detections, and by the multi-valued nature of the data for each image time as scatterers move in and out of shadows in the radar beam illumination.

3. CONTROL STRUCTURE: BLACKBOARD ARCHITECTURE

In the problem domain of signal interpretation, traditional approaches have avoided dealing with control issues. The problem-solving process is typically treated as a sequential set of fixed signal operations, numerical in nature. The order of events as well as the computational model are prescribed at the onset. The data passes from the input state to the output state through the same processing path irrespective of the data conditions, of any previous processing, or of the quality of the output. Such a problem-solving model is rigid and is useful only for the particular application for which it was designed. In addition, computational models make it difficult to incorporate and accumulate specific domain knowledge that we as humans readily bring to bear on the problem.

Developments in the field of machine intelligence have produced alternative models of the problem-solving process. Such models organize reasoning steps and relevant domain knowledge to produce a flexible means of controlling the events that lead to a satisfactory solution [2]. Current knowledge-based systems make use of a variety of problem-solving models, the most well known being backward (or forward) reasoning used extensively in production systems. The control works backward (or forward) from a goal (or initial) state reasoning to the desired final state. The path of the solution is dependent on a variety of conditions, subsolutions, and known facts about the current domain.

In addition to backward and forward reasoning approaches to the problem-solving process, certain applications have demanded less directed forms of control where the reasoning may involve a combination of backward and forward inferencing. Such control is an opportunistic reasoning model whereby action occurs when needed depending on the state of the domain knowledge and the partial solutions to the problem. Specific information processing is performed only at "opportune" times.

The blackboard architecture is a specific structured form of opportunistic reasoning. In general, as shown in Figure 3-1, the overall structure can be broken down into three main conceptual and physical entities: the blackboard, the knowledge sources, and the control mechanism. The blackboard is a structure which publically displays the relevant information that is known *a priori* or is accumulated during the problem solving process. The blackboard is essentially a database containing problem-specific data and the solution state. The organization of this database is typically a hierarchy of subsolutions corresponding to the natural domain and solution space organization. The knowledge sources are independent information processing units that incorporate problem-specific expertise and communicate directly with the blackboard. The knowledge source composition and intent varies widely since no structure is imposed on its contents. Knowledge sources are activated opportunistically, typically by matching the knowledge source requirement pattern against the current state of the blackboard. When a knowledge source has finished its processing, appropriate results are posted back to the blackboard, making them accessible to all other knowledge sources. A control mechanism mediates the actions of the blackboard and the knowledge sources. The controller activates all possible knowledge sources by matching all

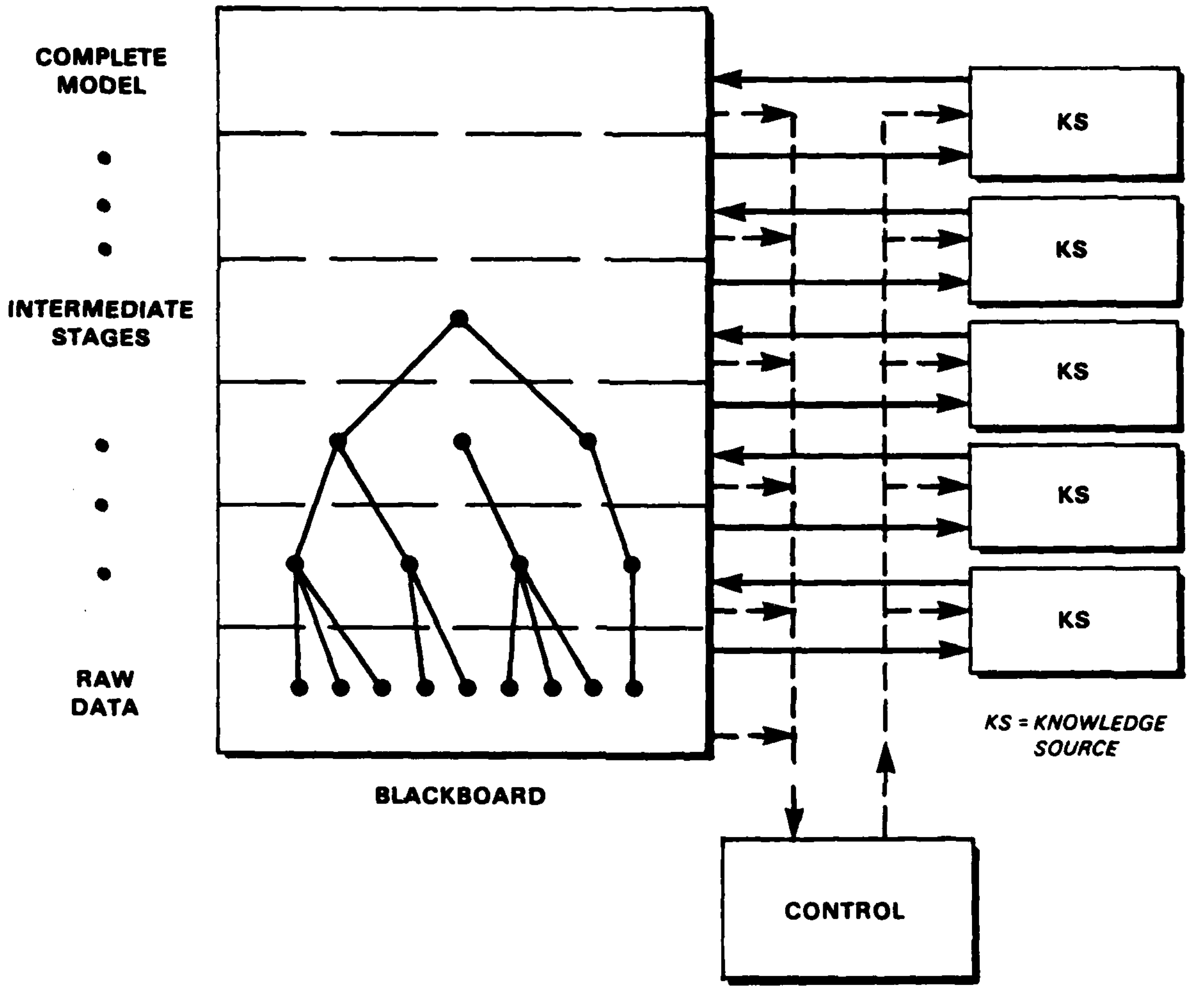

Figure 3-1. Blackboard Processing Architecture.

available knowledge sources against the blackboard. When more than one knowledge source is activated, the controller determines which knowledge source to fire, executes that knowledge source, and then either keeps track of previously activated knowledge sources or refreshes the activation list and starts over. This process continues until no more knowledge sources can be activated. The conditions for knowledge source activation and execution are flexible and vary greatly among applications.

One of the first applications of the blackboard model was on the Hearsay speech understanding system [4]. The domain of speech understanding breaks down into a natural hierarchical representation as the knowledge is distributed from the parametric, segmental, and acoustic levels up through lexical, phrasal, etc., levels. Multiple, diverse kinds of specific expertise are needed to incrementally solve this expansive problem. The blackboard architecture allows speech-specific expertise to be embedded in knowledge sources that communicate with a structure emulating the

natural domain organization of the speech understanding problem. In a manner similar to the speech domain problem, radar object modeling also requires multiple areas of expertise distributed through several layers of abstraction in order to derive a physical model and a motion solution from range-Doppler images. This observation has motivated our use of a similar architecture.

In the following sections, we will discuss our blackboard application in more detail. Our focus will be on the domain of radar object modeling and the demands it places on the problem solving strategy and architecture.

3.1 THE BLACKBOARD STRUCTURE

The blackboard structure currently in ROME is organized into domains that correspond to subsolutions of the radar object modeling problem as shown in Figure 3-2. Evidence is accrued into domains ranging from the lowest-level range-Doppler features and selected data segments up through the symbolic representations of the object and its motion and the catalog matching results. Each domain is further subdivided into types that distinguish classes within the domains such as spin estimates derived via different algorithms. When pieces of information on the blackboard are used by other knowledge sources, they are tagged as such to distinguish new evidence from previously-derived information. This allows opportunistic refinements during the problem solving process. The design of the blackboard does not dictate exclusively data-driven or model-driven processing, since both strategies can be exercised in the same framework.

3.2 THE KNOWLEDGE SOURCE STRUCTURE

The knowledge sources are data processing structures that contain all the information necessary for independently communicating with the blackboard and the controller and for its specific information processing. The knowledge sources are directed to take information from and return results to particular blackboard domains and subdomains. Patterns associated with individual knowledge sources are matched against the current patterns on the blackboard. If there is a match, the activation level for the knowledge source is set and the knowledge source is added to the controller's list of active knowledge sources to be scheduled or reset.

The majority of the knowledge sources point to algorithms which extract subsets of the data, identify and parameterize spinning and slipping scatterers on the spinning and precessing rigid body, and assess the fit of the derived features to the data. They encompass digital signal processing operations such as LMS estimation, spectral analysis, linear and nonlinear filtering, data transformations, clustering, etc. As different knowledge sources attempt to extract the same information using different techniques, additional knowledge sources serve to resolve redundant or conflicting information. Our physical knowledge is incorporated as a rigid body motion solution requires a unique spin rate for all scatterers. Additional knowledge sources build symbolic models from the derived features and match those against catalog models stored in the identical symbolic representation. The knowledge sources are independent and modular, thereby facilitating

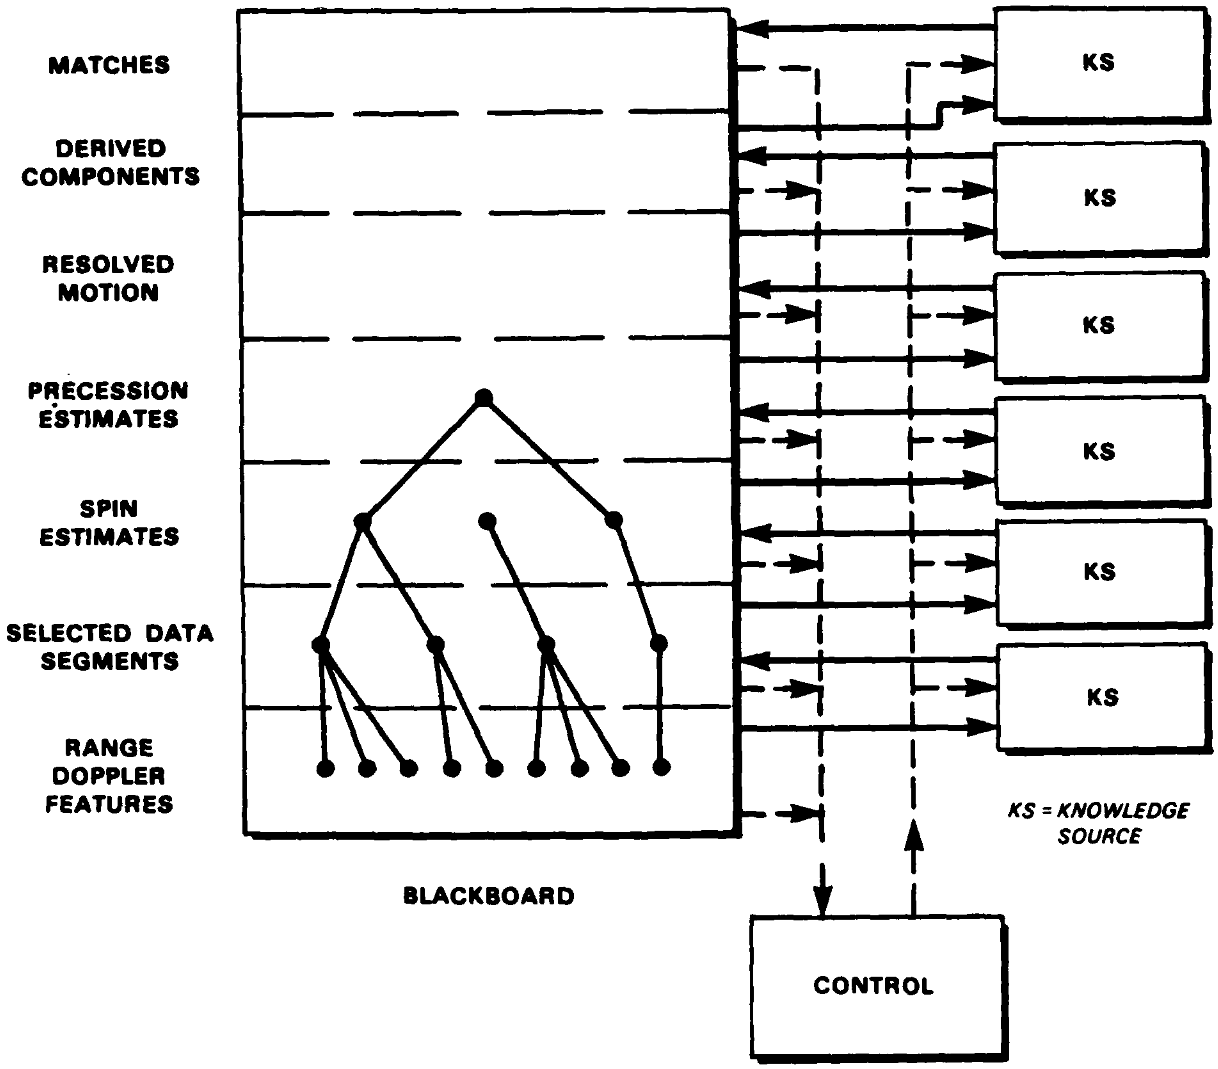

Figure 3-2. ROME Blackboard Processing Architecture.

the work of multiple algorithm developers within the project, multiple approaches to the same problem, and incremental changes and enhancements to the system.

3.3 THE CONTROL MECHANISM

The control mechanism dictates the sequence of operations. The controller first finds all possible active knowledge sources by matching the individual patterns within the knowledge sources against those on the blackboard. These patterns consist of requirements that must be true about the blackboard as well as conditions that must be met. When all possible knowledge sources are activated, the controller calls on the scheduler to determine which single knowledge source should be executed. The scheduler examines the scheduling syntax within all knowledge sources and orders them accordingly. Perhaps the simplest scheduling mechanism is the preassignment of fixed numerical priorities to each knowledge source. An alternative, which we have currently implemented, is domain-dependent, specifying that knowledge sources of one domain should execute before or after those of another domain. Eventually it might be appropriate to install dynamic scheduling where the schedule information for a knowledge source depended on the data, the blackboard, etc.

After scheduling a single knowledge source, the controller executes that knowledge source. All of the knowledge source inputs are obtained through the pattern matching against the blackboard. Remaining knowledge sources are ignored at this time. After the current knowledge source has executed and posted its results to the blackboard, all knowledge source activation levels are reset and the process is repeated starting from pattern matching all knowledge sources to determine the activation list.

This control mechanism accommodates data-driven and model-driven processing. Its flexibility stems from the modularity of the knowledge sources, the generic pattern matching and pattern specifications, and the ability of the knowledge sources to communicate independently with the domain-structured blackboard. We take advantage of this capability as initial attempts to match derived models against the catalog can cue lower level algorithms to look for more specific features or data segments.

4. FEATURE EXTRACTION

The feature extraction algorithms are imbedded in the knowledge sources as discussed previously. Their primary function is to extract scattering centers which evidence body spin and to extract slipping scatterers which exhibit precessing motion. The role of the feature extractors is also to select subportions of the data that best exhibit these phenomena. In addition, there are routines to resolve the results of redundant or conflicting knowledge sources. In this chapter, we will discuss how radar data is processed and handed to ROME for feature extraction and present detailed descriptions of the feature extraction algorithms that reside in the blackboard knowledge sources.

4.1 INITIAL DATA PROCESSING

The wideband radar returns are preprocessed and compensated for gross motion to produce images in range-Doppler space. Feature vectors extracted from images consist of local range bin maxima and the corresponding Doppler neighbors. This is the format in which we initially receive the data. From these local feature vectors, on a per image basis, scattering center peaks are located in range-Doppler space and tagged with their range, Doppler, radar cross section (RCS), polarization, and phase values. When all images have peak scattering centers marked, a data track is formed as the input to the ROME recognition system. This track is presented to the blackboard as the original set of range-Doppler features.

An example simulated data track is displayed in Figure 4-1, shown as Doppler vs. time with range, RCS, polarization, and phase suppressed. There is little noise in this particular data

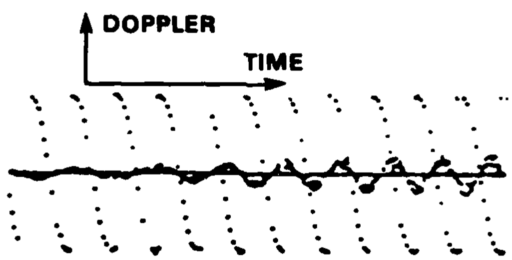

Figure 4-1. Typical Data Segment.

segment but the data does suffer from missing points and occlusion making the dimensionality of the data time-varying. Note the periodic structure within the track due to the returns from point scatterers rotating on the rigid body at the body spin rate. In this data, it is relatively simple for a human observer to extract the data points associated with the rotating point scatterers as well as the number and relative positions of the scatterers. The occlusion of the scatterers as they rotate behind the body and the number of scatterers do not pose an interpretation problem for a human observer who can resolve any ambiguities. However, these effects can cause irrecoverable

problems for classical signal processing methods. Note, for example, in the data the low Doppler returns from ring scatterers, such as manufacturing joins, exhibiting the precessing motion of the rigid body; these scatterers and their motion must be distinguished from the point scatterers. These difficulties reinforce our inclination to employ a knowledge-based approach in our algorithm design, control structure, and model representation strategies.

4.2 DATA CLUSTERING

The most challenging task in this modeling problem is to establish a correspondence between a set of points in range-Doppler-time space and the corresponding physical feature in the model space. Once points are assigned to the track of a given feature, extracting a parametric description is relatively straightforward. Thus one low-level knowledge source is a clustering algorithm, which attempts to partition the data space on the basis of range. An alternative approach is based on a smoothed histogram of RCS values in range. Each point is modeled as an impulse in range and then convolved with a smoothing function. Range partitions are determined by mimima or zero-crossings in a thresholded version of the resulting function. The individual groups are now assumed to contain either noise, tracks of spinning scatterers, or tracks of ring scatterers. Of course, a group may contain multiple tracks, and a track may be split across two groups; the subsequent processing must be robust enough to handle these cases. The partitioned data tracks are reported to the selected data segments domain of the blackboard, and tagged with type names corresponding to the type of further processing for which each track is most applicable.

4.3 SPIN EXTRACTION BY LPC FIT

In the data illustrated in Figure 4-1, the spinning scatterer tracks cover a wider Doppler spread than do the precessing ring scatterers; thus the spin extraction knowledge sources are next called to process the output of the clustering knowledge source (which has been posted on the blackboard). Three techniques are applied independently; all model, in a sense, the feature-extraction process applied by a human analyst. The first approach, discussed in this section, searches for peaks in Doppler values, indicating the presence of a spinning scatterer track. Modeling these tracks as sinusoids in Doppler-time space, we begin with three points near the peak to estimate frequency using a 1-parameter linear predictive coding (LPC) fit. The track is then extrapolated using a sinusoid model to locate successive points; as each successive point is added to the track, the LPC fit is repeated to iteratively refine the frequency estimate. The track search is terminated when we reach the break caused by occlusion of the spinning scatterer. With a frequency estimate based on all points of the track segment, we now correlate against a sine and cosine at that frequency to estimate the amplitude and phase of the best-fit sinusoid.

Moving one (estimated) period along the data, we now locate points which are hypothesized to belong to the same track. The LPC fit is repeated independently of the first estimate, and the two sinusoid parameter estimates are compared. If they are sufficiently close, the frequency estimate is refined using the track zero-crossings, and the fit of amplitude and phase is repeated

using the combined data from both segments. This process is repeated to the limits of the data segment.

This process is termed "bottom-up" or "data-driven"—given the data points, we extract parameters according to some fit criterion. We now take a "top-down" viewpoint: expecting a sinusoidal pattern, we compare the parametric fit with the data, identifying those points which can be explained with the given model component. Since there will, in general, be more than one scattering center track in a given data cluster, we must repeat the search procedure, successively extracting tracks and interpreting them as spinning scatterers on the model. This stage of the processing ends when no further tracks can be extracted; the remaining data points are assumed to be anomalies (poor fits or noise) or the tracks of ring scatterers, which do not exhibit periodic occlusion.

4.4 SPIN EXTRACTION USING THE HOUGH TRANSFORM

An alternative approach is based on the observation that the data of Figure 4-1 includes a set of quasi-linear segments. This suggests that the Hough transform [5] may be employed to map from time-Doppler to slope-intercept space, where the line segments may be detected as local peaks. An interesting problem arises in the interpretation of the Hough transform peaks: although we are interested in clustered points along regularly spaced lines with negative slope, which the eye readily extracts from the data of Figure 4-1, a simple peak detection in the Hough slope-intercept plane leads to the detection of many lines present in the data, but inconsistent with our physical understanding. Here is where the machine intelligence component of the processing becomes evident: even within the knowledge sources, we must often apply a significant amount of reasoning in using the output of the signal processing algorithms. Thus a logical filter must be applied to the Hough line detections on the basis of the time span of their associated points and on the magnitude and sign of their slope. Figure 4-2 displays the selected Hough line detections superimposed on the same segment of data as seen in Figure 4-1.

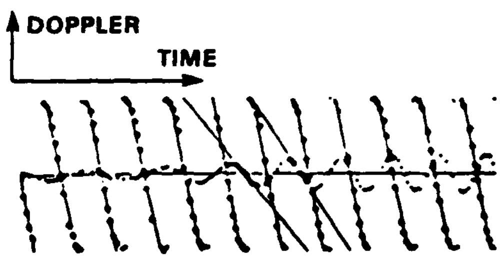

Figure 4-2. Hough Transform Line Detections.

An estimate of spin frequency may be made by observing that typically 40-50% of a period is spanned by a line segment. Alternatively, an LPC frequency estimate may be applied to the point cluster for each line. The number of spinning scatterers is then extracted by an analysis

of periodicities in the zero-crossings for a given spin frequency estimate, allowing for missed line detections and for anomalous detections representing false lines, as seen in Figure 4-2. The primary test in this portion of the modeling is one of consistency among the spin periods observed, according to the implicit rigid body assumption.

4.5 DOPPLER-TIME CLUSTERING

A third approach is warranted in cases where the Doppler coordinates of the data points are too noisy to link with the LPC approach and the presence of multiple scatterers leads to ambiguous detections with the Hough transform line extraction. The Doppler-time clustering algorithm begins with the evaluation of a track autocorrelation to detect periodic structure in the data. An estimate is made of the fundamental body spin period (e.g., from the highest autocorrelation peak or as evidenced by the amplitude spectrum of the track autocorrelation). If this estimate is ambiguous, several alternative values of spin period might have to be carried in parallel through succeeding processing stages before a fit estimate can resolve the uncertainty. The track data is then wound up on the selected spin period T, with points at time $t + nT$ overlaid on those at time t.

A spreading function $w(t, D)$ is then used to smooth the points in time and Doppler, yielding the function

$$A(t, D) = \sum_i \sigma_i w(t - t_i, D - D_i) \tag{4.1}$$

where σ_i and (t_i, D_i) are the amplitude and (time, Doppler) coordinates of the ith track point, and the sum is taken over all points on the wound-up track. The function $A(.,.)$ is evaluated on a fine grid in the time-Doppler plane, local peaks are identified, and (t, D) points are collected along ridges emanating from each peak. A sinusoid of period T is fit to the points forming each ridge. This set of sinusoids is now compared with the original track data, and selected in the order of best fit to data until a lower fit quality threshold is reached. Estimates of sinusoid amplitude and phase are refined through comparison with close data points. It is at this point that ambiguities in spin rate can be resolved to obtain the appropriate submultiple of the apparent repetition frequency, also indicating the true number of discrete scatterers.

4.6 RESOLUTION OF SPIN EXTRACTION

The last three sections have outlined three different approaches to extracting the same information from the data signature. The choice of fitting algorithm may be decided on an *a priori* basis, depending on the data or object type. Alternatively, the three algorithms may be run in parallel on the same data, and an *a posteriori* selection made on the basis of fit quality. The spin fit resolver could then be viewed as a fourth knowledge source. The results of the various spin extraction algorithms are posted to the appropriate types within the spin estimates domain of the blackboard. The spin fit resolver takes those pieces of evidence and reports a single evaluated

result to the resolved motion domain of the blackboard. The flexibility of this modular structure has expedited the construction of our system, with algorithmic implementations by different analysts merged at well-defined interfaces.

4.7 PRECESSION ESTIMATION

Since precession is variable in amplitude and frequency, especially in the presence of forces induced by atmospheric interactions, an appropriate model must be chosen for its description. Letting ψ be the precession magnitude, θ the precession angle, d the distance from ring scatterer to reference point, and α the aspect angle as before, we have a Doppler signature of (approximately) $f_D(t) \approx \psi(t)\dot{\theta}(t)d\cos[\theta(t)]$. Here, $\psi(t)$ and $\theta(t)$ are smooth functions of time (we chose low-order polynomials over disjoint time intervals). The DSP operations employed include sign extraction from a weighted sum of Doppler values for each time, median filtering to remove dropouts, zero-crossing detection, sinusoidal fit over consecutive (estimated) half-periods, and an LMS polynomial fit of magnitude and net angle for successive time segments. In the design of the fit algorithm for this time-varying phenomenon, a balance must be struck between the closeness of the match and the length of the time segment over which it is valid. The results of the precession estimation algorithm are reported to the precession estimates domain of the blackboard. As with spin extraction, a precession resolver knowledge source examines the results for ambiguities or conflicts and posts final results to the resolved motion domain.

5. MODEL REPRESENTATION AND MODEL MATCHING

The choices of a model representation and a model matching strategy are tightly coupled and depend heavily on the objects being modeled, on the nature of the data, on the types of features extractable from the data, and on the general approach to the problem (e.g., data-driven versus model-driven). The characteristics of range-Doppler images of fairly simple rigid objects with a tractable motion solution provide several specifications for a model representation. First, we can image the entire object through several rotations since its spin rate is considerably lower than the imaging rate and our observation time can be on the order of several periods of rotation. Thus, a 3-dimensional description of the located scatterers can be derived. Second, the distinguishing factor between different imaged objects is the general body structure and motion solution, as opposed to more detailed numerical specifications. Therefore, a suitable object model should symbolically represent the structural components, their interrelationships, and pointers to motion solutions and numerical quantifiers. Finally, while there is a set of known objects for which we would like to build and store models to be compared with imaged objects, we also need to acquire new models from the data as new previously uncataloged objects are presented to the system. Thus, it is important to be able to derive a workable model directly from the data. The data-derived model should be identical in form to the stored catalog models, thus facilitating building and generalizing new catalog entries directly from the data.

5.1 MODEL REPRESENTATION

The model representation currently in ROME is a simple version of a semantic network. A semantic network is a network of nodes and links, with nodes typically corresponding to objects or concepts and links symbolizing relations between the nodes [6],[7]. In addition, unary property nodes can point to and quantify either nodes or links. This is a fairly common model representation for structured objects or scenes from imagery [8], [9].

The semantic network derived from images of simple objects consists of scatterers, the relationships between them, their motion, and quantifiers associated with them as shown by the example network in Figure 5-1.

Nodes in the network correspond to point scatterers, ring scatterers, extended returns, and the entire rigid body. The relational links between the nodes represent relative differences in range or angle. These are differential relationships, not absolute quantifiers. Quantifiers pointing to the nodes contain more detailed positional information such as crossrange, radar cross section, or range. The motion, particularly the spin rate, is a property attached to the entire body node, as all scatterers extending from the body node must rotate at that spin rate. Thus, the scatterer nodes inherit the motion solution from the rigid body node. An additional knowledge source takes the current derived components and creates a semantic network. The data-derived features are thus in the same format as the catalog of objects available for comparison.

Due to the relatively simple structure of objects such as reentry vehicles, the corresponding semantic network representation is simple with few nodes, connections, and layers of abstraction.

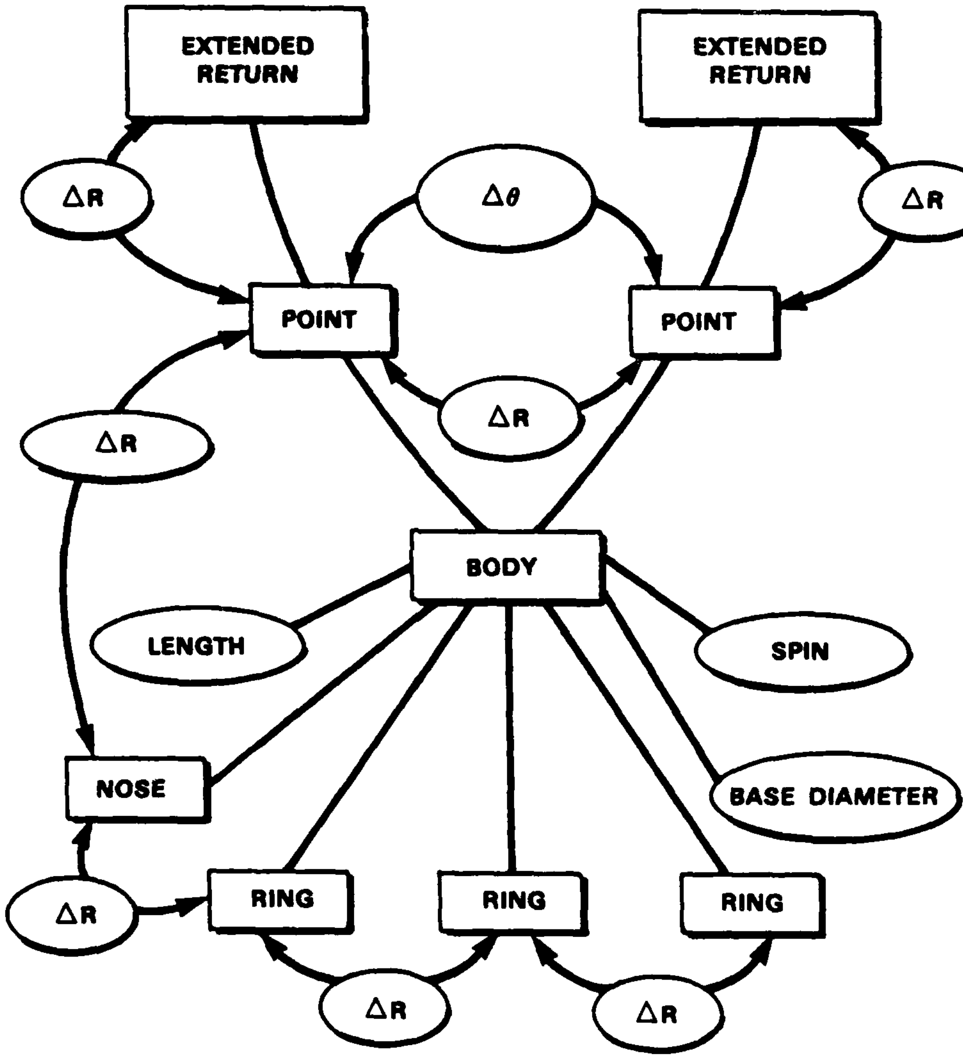

Figure 5-1. Rigid Body Semantic Network Model.

As the system is exposed to more complex objects such as satellites, the semantic network will provide a powerful means of representing and manipulating these complex structures, as well.

5.2 MODEL MATCHING

Once a data-derived semantic model exists, it can be matched against a catalog of semantic models. The catalog models are established directly from known physical models or from data-derived models that were interesting or different from current catalog models.

The matching strategy consists of finding the best correspondence between parts on the derived model and the catalog model. This correspondence is evaluated and compared to the correspondences from the other catalog models. Due to the symmetry of the objects and the occasional discrepancy in number of parts between the derived model and the catalog model, there are usually several potential correspondences between point and ring scatterers. These correspondences are based on the general structure rather than on quantifiers. The correspondences are constrained by the fact that range and angular position information limits the permutations possible for these correspondences.

Once a set of realizable correspondences is established between the point/ring scatterers from the derived model and the point/ring scatterers from the catalog model, these correspondences must be evaluated and scored to assess the best fit. The basis for evaluation is finding the correspondence that minimizes the distances in property space. For example, assume in a certain correspondence that a ring from the derived model is associated with a ring from the catalog model. The properties attached to both rings (such as range and RCS) are plotted in their property space and Euclidean distances are computed between derived and cataloged properties. This is done for all parts in the correspondence, and for all correspondences. The correspondence with the greatest number of minimum distances across its parts is marked as best. The scoring metric for each correspondence is the number of minimizations in property space less the number of missing parts between the derived and catalog models. Some cost must be associated with a derived model having either too few or too many parts compared to the catalog model. The resulting normalized scores are compared across the catalog, and the maximum chosen as the best match from the catalog to the derived model. Note that the spin rate and the apparent number of point scatterers are mutually dependent pieces of evidence, which may be used to refine an object model in a model-driven processing pass.

5.3 GENERALIZATION

Due to the variety of missions of the radars from which we obtain data, images of unknown objects are often presented to the system. In these cases, the data corresponds to objects for which no *a priori* models exist. An important part of our system is the ability to generalize a semantic model from any number of data segments associated with a particular object.

Generalization is a subproblem in the more general field of learning [10]. Winston [11] advocates an inductive learning strategy known as learning by example. An external teacher presents

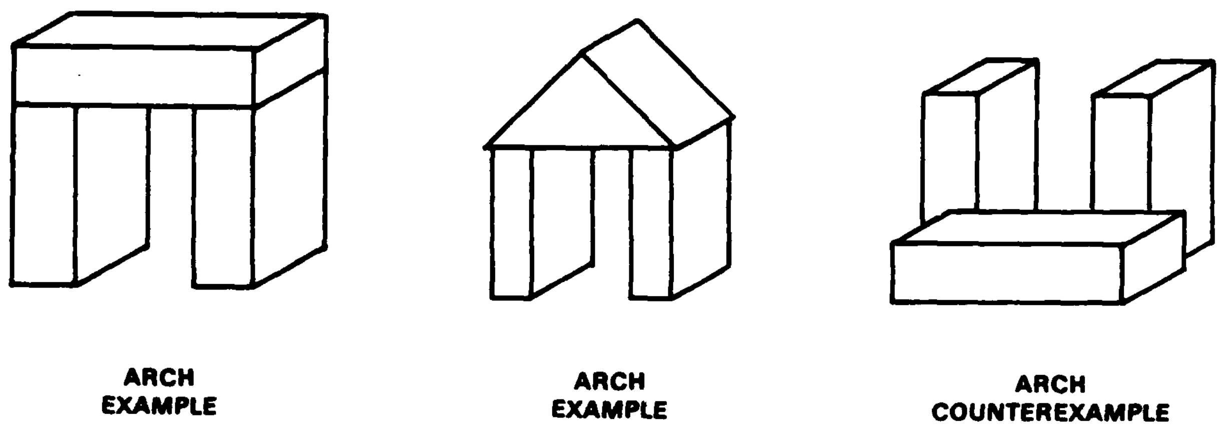

Figure 5-2. Winston's Learning the Structure of an Arch by Example.

the system with examples and counterexamples and the system induces the general concept description. Winston's classic analogy, as seen in Figure 5-2, is that of learning the structure of an arch. The system is presented with the first two configurations as positive examples of what an arch should look like. The third case is presented as a counterexample or how an arch should *not* appear. From these examples, the system induces the general concept of an arch and builds an internal representation of that concept. As new inputs are presented to the system, they are matched against the current models and recognized as an example of an arch or not. New inputs can also be used to continually refine the internal representations.

As opposed to a true learning system that continually changes its knowledge and performance based on examples, we approach generalization chiefly as a way to train the system as to what a particular uncataloged object should or should not look like. This information is then entered into the current catalog. Once the training is complete, the system performs in its normal mode, treating the generalized model as a catalog entry.

The generalization process is guided by a human operator who selects acceptable data-derived models for a given object as well as counterexamples. Acceptable models as well as counterexamples are determined based on the expertise of the human operator. The resulting models can thus be ambiguous due to the complexity of the data and the inconsistency among experts. These ambiguities are to be resolved through generalization. Two types of models are generalized through the same procedure: example models and counterexamples models. All the data-derived models selected are the results of the initial processing of the ROME system (up to the matching stage) to produce semantic models just as discussed previously. The first derived model initializes the generalized model. As other examples are presented, the generalization is applied to associate new structural components with current structural components in the generalized model. When an association is made (similar to the matching correspondence problem), the structural components increases a number-of-occurrences counter. Each structural component keeps track of the number of associated occurrences as well as the total number of training instances. Thus, each component in the generalized model has a weight attached to it equal to the number of occurrences divided by the number of training sets. This weight, which is between 0 and 1, is used to estimate the importance of that structural component in the matching. Components that are not associated with existing parts in the generalized model are added as additional parts, but their weight is low compared to parts that do find appropriate associations.

Once an example model and counterexample models are generalized for a particular object, both types of models are entered as the catalog entry for the unknown object, and the system performs its normal operations. As new data comes into the recognition system and a semantic model is built, the derived model is compared against both example and counterexample models. A possible match is considered when the derived model matches the example model better than it matches the counterexample model. The number of counterexamples needed for a particular object depends on the diversity of the models in the entire catalog and the inherent variability in the data and in the data processing.

Initial results from the generalization of models have been encouraging. Inappropriate data-derived models are usually discarded, since they correspond more strongly to the counterexample

models. Exceedingly well-derived models clearly stand out, since the margin between the example model match and counterexample model matches is significant. A point of sensitivity in the generalization scheme is careful selection of distinguishable counterexamples for generalization.

6. PERFORMANCE EVALUATION

The ROME object modeling system has been evaluated on several sets of data including images of objects with known dimensions as well as those with unknown characteristics. We have two cataloged known objects, referred to here as *Object A* and *Object B*, as well as data for several unknown objects. The appearance of the data varies from clean, simple, fixed scatterer returns to noisy, sparse, flexible scatterer returns, as suggested by the simulated data segments shown in Figure 6-1.

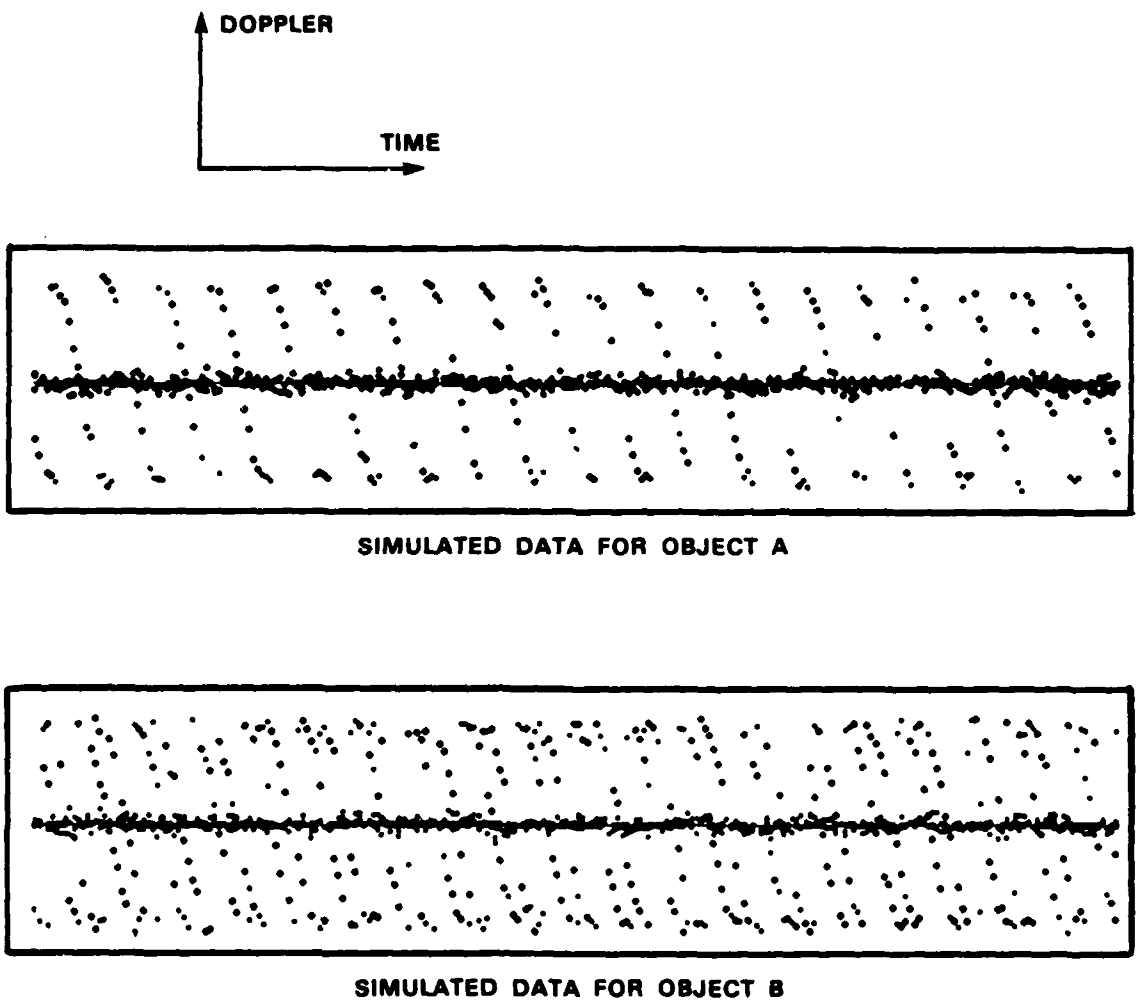

Figure 6-1. Simulated Data Examples.

In general, the results from testing the ROME system on real data have been encouraging. The models produced agree with manual analysis in the case of unknown objects, and with known

object dimensions in the case of *Objects A* and *B*. In addition, the model matching procedure successfully associates the derived model with the correct model from the catalog despite anomalies in the derived model. The addition of top-down processing (i.e., cueing the feature extractors with high-level information from potential catalog matches) further enhances the matching results in cases where bad estimates of spin motion lead to erroneous models in terms of number of spinning scatterers and crossrange dimensions. For the remainder of the evaluation, our focus will be on the objects with known dimensions and motion. This provides a true reference from which quantitative results can be assessed.

The objects with known dimensions, *Object A* and *Object B*, have nearly identical physical dimensions and motion solution, with the exception of two additional point scatterers on *Object B*. This similarity increases the difficulty in distinguishing these objects from each other. However, the system's ability to make this distinction gives us reassurance that the system will perform well in foreseeable applications where a greater number of less correlated objects are to be classified.

The images for *Objects A* and *B* were subdivided into 24 and 10 test tracks, respectively. The test tracks vary from one to four seconds (50 to 200 images) in duration. The catalog used in the evaluation consisted of models for *Objects A* and *B* as well as models for 8 similar synthetic objects. The synthetic models exhibit a realistic range of point and ring scatterer structures and motion solutions. The system results are shown in Figure 6-2. The catalog model of *Object A*

	CLASSIFIED AS		
OBJECT	A	B	OTHER
A (24)	63%	21%	16%
B (10)	0%	80%	20%

Figure 6-2. ROME Evaluation Results.

was correctly matched to the data-derived model 63% of the time. Misclassifications of data from *Object A* were mostly due to errors in spin rate estimation, resulting in overestimates of the number of spinning scatterers and the derived model appearing more like *Object B*. Classification of *Object A* as one of the synthetic models stemmed from noisy data segments or data exhibiting large precession motion interfering with spin and scatterer extraction. The data-derived models for *Object B* were correctly classified 80% of the time. Interestingly enough, none of *Object B*'s data was classified as *Object A*, perhaps due to the less structured nature of this data causing overestimates in the number of spinning scatterers. Another point of interest in the matching results is that the scores returned with each match accurately reflect the quality of match. This gives us confidence that the scoring metric is meaningful in relation to distinguishing the objects by their physical structure. Also, based on the returned match scores, it is possible to set a threshold on acceptable matches or to determine the appropriateness of further processing.

It should be emphasized at this point that the preceding results represent an initial evaluation with limited data. The availability of data has been a limiting factor in development and evaluation, but remedies to this shortage of data are forthcoming in the near future. At that time, the system will be fully exercised and newer enhancements, such as top-down processing and model generalization, will be thoroughly tested.

7. SUMMARY

In this report, we have presented a unique approach to solving a radar object modeling problem. The appropriateness of machine intelligence in this endeavor stems from three sources. First, in the radar domain, manual analysis and interpretation of the imagery is straightforward and readily learned. Second, machine intelligence methods are less sensitive to the condition and amount of data. Finally, machine intelligence offers flexible representations and control structures in which to imbed domain-specific techniques.

The overall goals of our project are, given range-Doppler images of a simple rotating object, to extract primitive and high-level features from the images, formulate a rigid body model, and identify the object with an interpretation consistent with that of a human analyst. Our efforts have demonstrated the capabilities to extract and identify physical models from the radar images through the use of appropriate signal processing primitives, knowledge representations, and control structures. The machine intelligence approach enables classification of models with no *a priori* statistical information, with minimal actual model expectations, and with uncertainties in the data and parameter bounds. The power lies in the data-driven fashion in which signal processing algorithms are directed and the ability to match ambiguous and partial derived models at the symbolic level. The semantic model representation incorporates both static and dynamic information. This representation is able to cope with a lack of information and lends itself to generalization of unknown objects. The "semantics" of the 3-dimensional object are fully imbedded in the representation and not interpreted by an external program. The blackboard control philosophy allows for modular, task-specific expertise to be accessed opportunistically. The problem solution is thus acquired incrementally as evidence is accumulated, implementing a bootstrapping problem solving strategy.

As the development of the ROME system continues, we intend to acquire a larger database of imaged objects in order to fully exercise the system's capabilities. Initial evaluation results are encouraging with a limited variety of data and object models. As more experience is gained with generalization and model-driven processing as well, the system's ability to deal with a broad range of data conditions and uncataloged objects will be strengthened.

The ROME system is part of the ongoing activities in the Machine Intelligence group. Several people have been involved in its development; among them, Don H. Johnson and Ellen M. Glover contributed significantly to the architecture and algorithm design.

REFERENCES

1. G.E. Kopec, A.V. Oppenheim and R. Davis, "Knowledge-Based Signal Processing," *Trends Perspectives in Signal Process.* 2(2):1–6 (July 1982).

2. H. Penny Nii, "Blackboard Systems," *AI Magazine*, Vol. 7 (3 and 4) (Summer 1986).

3. D. A. Ausherman, A. Kozma, J. L. Walker, H. M. Jones, and E. C. Poggio, "Developments in Radar Imaging," *IEEE Trans. Aerosp. Electron. Syst.* AES-20:363–400 (July 1984).

4. V. R. Lesser, R.D. Fennell, L. D. Erman, and D. R. Reddy, "Organization of the HEARSAY-II Speech Understanding System," *IEEE Trans. Acoust. Speech, Signal Process.* ASSP-23(1) (February 1975).

5. D. H. Ballard and C. M. Brown, *Computer Vision* (Prentice-Hall, Englewood Cliffs, NJ, 1982).

6. A. Barr and E.A. Feigenbaum, editors, *The Handbook of Artificial Intelligence.* Volume 1 (William Kaufmann, Inc., Los Altos, CA, 1981).

7. P.H. Winston, *Artificial Intelligence* (Addison-Wesley, Reading, MA, 1984).

8. J.H. Connell, "Learning Shape Descriptions: Generating and Generalizing Models of Visual Objects," Technical Report 853, MIT Artificial Intelligence Laboratory (September 1985).

9. T. E. Weymouth, "Using Object Descriptions in a Schema Network for Machine Vision," Technical Report 86-24, University of Massachusetts, Computer and Information Science (1986).

10. R.S. Michalski, J.C. Carbonell, and T.M. Mitchell, *Machine Learning* (Morgan Kaufmann, Inc., Los Altos, CA, 1983).

11. P.H. Winston, "Learning Structural Descriptions from Examples," Technical Report 231, MIT Artificial Intelligence Laboratory (1970).

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY 2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
4. PERFORMING ORGANIZATION REPORT NUMBER(S) Technical Report 818	5. MONITORING ORGANIZATION REPORT NUMBER(S) ESD-TR-88-195
6a. NAME OF PERFORMING ORGANIZATION Lincoln Laboratory, MIT 6b. OFFICE SYMBOL *(If applicable)*	7a. NAME OF MONITORING ORGANIZATION Electronic Systems Division
6c. ADDRESS *(City, State, and Zip Code)* P.O. Box 73 Lexington, MA 02173-0073	7b. ADDRESS *(City, State, and Zip Code)* Hanscom AFB, MA 01731-5000
8a. NAME OF FUNDING/SPONSORING ORGANIZATION U.S. Army Strategic Defense Command — Huntsville 8b. OFFICE SYMBOL *(If applicable)* CSSD-H-YD	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F19628-85-C-0002
8c. ADDRESS *(City, State, and Zip Code)* Sensors Directorate P.O Box 1500 Huntsville, AL 35807-3801	10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO.: 62322C PROJECT NO.: 342 TASK NO.: WORK UNIT ACCESSION NO.:

11. TITLE *(Include Security Classification)*
Machine Intelligence Applied to Radar Object Modeling

12. PERSONAL AUTHOR(S)
Ann Marie Aull and Robert A. Gabel

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT *(Year, Month, Day)*	15. PAGE COUNT
Technical Report	FROM ______ TO ______	1988, October 12	42

16. SUPPLEMENTARY NOTATION
None

17. COSATI CODES

FIELD	GROUP	SUB-GROUP

18. SUBJECT TERMS *(Continue on reverse if necessary and identify by block number)*

artificial intelligence
machine intelligence
image understanding
blackboard architecture
object recognition
radar signal processing

19. ABSTRACT *(Continue on reverse if necessary and identify by block number)*

In this report, we discuss a machine intelligence approach to modeling simple space objects from radar range-Doppler images. The data is multidimensional in nature with additive noise, distortion, and missing points. The relevant features to be extracted from the radar data include the position of all scattering centers in body coordinates, identification of the scattering center type (sphere, corner, edge, etc.), and motion parameters for the object. Our goal is to produce a representation of the imaged 3-dimensional object that is appropriate for recognizing the object as an example of something we have seen before and cataloged, recognizing the object as an uncataloged object, or determining discrepancies between the recognized object and our expectations of its appearance. We have built a recognition system with three major conceptual modules. The first of these is a set of signal processing primitives that are directed at the data to select subsets of data, extract features, and compare extracted features with the data to produce confidence measures. The second major module is the semantic model building and matching scheme. This component takes the data-derived features and produces a semantic model which is then matched against a catalog of stored semantic models for object identification. The final conceptual module in the system is the control structure which is based on a blackboard architecture from the field of machine intelligence. In this report, we will discuss the issues in range-Doppler imaging and describe the major system modules in more detail including motivations for our approach. We will also provide some performance evaluation results. Finally, summary and directions for future work will be presented.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified
22a. NAME OF RESPONSIBLE INDIVIDUAL Lt. Col. Hugh L. Southall, USAF	22b. TELEPHONE *(Include Area Code)* (617) 981-2330 22c. OFFICE SYMBOL ESD/TML